



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/496,844	02/02/2000	Patrick Knebel	10971393-1	6757

22879 7590 08/15/2003

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER

HUISMAN, DAVID J

ART UNIT	PAPER NUMBER
----------	--------------

2183

12

DATE MAILED: 08/15/2003

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No. 09/496,844	Applicant(s) KNEBEL ET AL.	
Examiner David J. Huisman	Art Unit 2183	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 01 July 2003.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1,3-5,7-15 and 18-22 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1,3-5,7-15 and 18-22 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☒ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
- 11) ☐ The proposed drawing correction filed on _____ is: a) ☐ approved b) ☐ disapproved by the Examiner.
If approved, corrected drawings are required in reply to this Office action.
- 12) ☐ The oath or declaration is objected to by the Examiner.

Priority under 35 U.S.C. §§ 119 and 120

- 13) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
* See the attached detailed Office action for a list of the certified copies not received.
- 14) ☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. § 119(e) (to a provisional application).
a) ☐ The translation of the foreign language provisional application has been received.
- 15) ☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. §§ 120 and/or 121.

Attachment(s)

- 1) ☐ Notice of References Cited (PTO-892) 4) ☐ Interview Summary (PTO-413) Paper No(s). _____
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948) 5) ☐ Notice of Informal Patent Application (PTO-152)
- 3) ☐ Information Disclosure Statement(s) (PTO-1449) Paper No(s) _____ 6) ☐ Other: _____

Art Unit: 2183

DETAILED ACTION

1. Claims 1, 3-5, 7-15, and 18-22 have been examined.

Papers Submitted

2. It is hereby acknowledged that the following papers have been received and placed of record in the file: #11. Amendment "C" as received on 7/1/2003.

Specification

3. The disclosure is objected to because of the following informalities: The copending application mentioned on page 9, lines 9-12, should no longer be listed as copending since it has been allowed. Also, its corresponding patent number should be provided.

Appropriate correction is required.

Claim Objections

4. Claims 10, 18, 21, and 22 are objected to because of the following informalities: For increased clarity, it is recommended that the applicant replace "SSE" with what it actually represents. Appropriate correction is required.

5. Claim 18 is dependent on a cancelled claim (17). Appropriate correction is required. For purposes of this examination, claim 18 will be viewed as being dependent on claim 12.

Art Unit: 2183

Claim Rejections - 35 USC § 103

6. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

7. Claims 1, 3, 4, 7-8, 12-15, and 19 are rejected under 35 U.S.C. 103(a) as being unpatentable over Col et al., U.S. Patent No. 6,330,657 B1 (as applied in the previous Office Action, mailed on 12/20/2002, and herein referred to as Col) in view of Shang et al., U.S. Patent No. 5,764,971 (herein referred to as Shang).

8. Referring to claim 1, Col has taught a method for processing software instructions comprising:

a) decomposing a macroinstruction into a plurality of microinstructions. See Fig.4, steps 402 and 404.

b) forcing the parallel issue of at least two of the plurality of microinstructions simultaneously. See column 3, lines 52-56.

c) executing the at least two microinstructions simultaneously, in lockstep using functional units in a floating-point unit. See column 3, lines 31-35, and Fig.6 (note in cycle 7 that two microinstructions are executed in parallel). Furthermore, note from column 20, lines 32-38 and note that this parallel execution can occur using multiple floating-point functional units.

d) Col has not explicitly taught:

d1) determining whether an exception occurs in any of the microinstructions, before writing results of the executing to result registers.

Art Unit: 2183

d2) if an exception occurs in any of the microinstructions, canceling all of the microinstructions and preventing the results of the executing from being written to the result registers.

d3) if no exception occurs in any of the microinstructions, writing the results of the executing to the result registers.

However, exceptions and the advantages of detecting exceptions are well known and expected in the art. In general, an exception is an interruption to the normal flow of program control, caused by the program itself or by executing an illegal instruction. Shang has taught a system in which macroinstructions are translated into a plurality of microinstructions and if an exception is detected within any one of those microinstructions, then the rest of the microinstructions are cancelled and results are not written to the result registers. Otherwise, if no exception is detected, then the results of the microinstructions are written to the result registers. See Fig.6 and note that if exceptions (interrupts) have been detected in at least one of the microinstructions at step 104, then the system is flushed (cancellation of each microinstruction) at step 108. If no exceptions were detected at step 104, then the results are written to the result registers at step 106. In Col's system, since a macroinstruction is broken up into microinstructions, an exception in a single microinstruction would mean that an exception has occurred in the overall macroinstruction, and therefore, all of the microinstructions that represent a single macroinstruction, should not be able to change the state of the system. An advantage of Shang's scheme would be to avoid having to undo the changes made by the undesired execution of a microinstruction that is part of a faulty macroinstruction. This will prevent a reduction in throughput in that the extra time required to perform an undo-operation would not be necessary.

Art Unit: 2183

Therefore, in order to maximize the efficiency of the overall system, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system such that it employs exception detection among microinstructions as taught by Shang.

9. Referring to claim 3, Col in view of Shang has taught a method as described in claim 1. Col has further taught that the microinstructions are executed on separate execution units, but appear as though they were executed on a single execution unit. See column 20, lines 32-38. Note that multiple floating-point execution units are used per clock cycle in order to execute multiple microinstructions. For example, the microinstructions of cycle 7 in Fig.6 must be executed on different execution units if they are executed in parallel. Then the results of each microinstruction are written to the appropriate storage via store logic (component 420 of Fig.4). See column 17, lines 3-16. Note that the store logic retrieves all of the results from each microinstruction execution and writes the data to the appropriate place. Finally, from Fig.6 (cycle 7), it should be realized that the separate, but parallel, execution of LD T1,[BX] and ADD AX,T1, will produce a result that is expected for the ADD AX,[BX] macroinstruction.

10. Referring to claim 4, Col in view of Shang has taught a method as described in claim 1. Col has further taught that all of the microinstructions are executed on the same clock cycle. See Fig.6, cycle 7, for instance.

11. Referring to claim 7, Col in view of Shang has taught a method as described in claim 1. Furthermore, Col has taught that the system allows a single instruction to operate on multiple single-precision floating-point values. See column 3, lines 61-63.

12. Referring to claim 8, Col in view of Shang has taught a method as described in claim 1. Col has not explicitly taught that a flag is updated based upon a result of the execution of the

Art Unit: 2183

microinstructions. However, it is well known in the art that processors contain a status register. The status register contains bits that are set or cleared based on the result of an operation. Some of the more common flags are ones that indicate a result of zero, a negative number, and overflow. These flags can then be checked in conditional situations, such as branches. Therefore, it would have been obvious to one of ordinary skill in the art to update a flag based upon the result of the execution of the microinstructions.

13. Referring to claim 12, Col has taught a computer system comprising:

a) a processor comprising:

1) a floating-point unit comprising a plurality of functional units adapted to execute microinstructions. See Fig.4, component 414, and column 20, lines 32-38.

2) Col has not explicitly taught a computer system with a ROM. However, Official Notice is taken that ROMs are well known and expected in the art. Processors contain Read-Only Memory to store essential software of the computer. Because it's non-volatile memory, ROM does not lose its contents when the power is turned off. Therefore, a ROM chip is used to store control programs for the computer, such as the bootstrap program (which tells the computer how to start and load the operating system) and other types of configuration information. As a result, it would have been obvious to one of ordinary skill in the art at the time of the invention to provide some type of Rom in Col's system.

3) a plurality of floating-point registers. See column 5, lines 49-52.

b) wherein the processor is configured to emulate an instruction set by:

Art Unit: 2183

- 1) decomposing a macroinstruction into a plurality of microinstructions. See Fig.4, steps 402 and 404.
- 2) forcing the parallel issue of at least two of the plurality microinstructions simultaneously to the functional units. See column 3, lines 52-56.
- 3) Col has not explicitly taught determining whether an exception occurs in any of the functional units, setting result registers for results of each of the functional units only if no exception occurs in any of the functional units, and if an exception occurs in any of the microinstructions, canceling all of the microinstructions and preventing the setting of result registers for all of the functional units. However, exceptions and the advantages of detecting exceptions are well known and expected in the art. In general, an exception is an interruption to the normal flow of program control, caused by the program itself or by executing an illegal instruction. Shang has taught a system in which macroinstructions are translated into a plurality of microinstructions and if an exception is detected within any one of those microinstructions, then the rest of the microinstructions are cancelled and results are not written to the result registers. Otherwise, if no exception is detected, then the results of the microinstructions are written to the result registers. See Fig.6 and note that if exceptions (interrupts) have been detected in at least one of the microinstructions at step 104, then the system is flushed (cancellation of each microinstruction) at step 108. If no exceptions were detected at step 104, then the results are written to the result registers at step 106. In Col's system, since a macroinstruction is broken up into microinstructions, an exception in a single microinstruction would mean that an exception has occurred in the overall macroinstruction, and therefore, all of the

Art Unit: 2183

microinstructions that represent a single macroinstruction, should not be able to change the state of the system. An advantage of Shang's scheme would be to avoid having to undo the changes made by the undesired execution of a microinstruction that is part of a faulty macroinstruction. This will prevent a reduction in throughput in that the extra time required to perform an undo-operation would not be necessary. Therefore, in order to maximize the efficiency of the overall system, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system such that it employs exception detection among microinstructions as taught by Shang.

14. Referring to claim 13, Col in view of Shang has taught a computer system as described in claim 12. Col has further taught that the processor is further configured to emulate the instruction set by executing all of the microinstructions. See column 3, lines 31-35, and Fig.6 (note in cycle 7 that two microinstructions are executed in parallel).

15. Referring to claim 14, Col in view of Shang has taught a computer system as described in claim 13. Furthermore, it has been noted that the computer system of claim 14 performs the method of claim 3. Therefore, claim 14 is rejected for the same reasons set forth in the rejection of claim 3 above.

16. Referring to claim 15, Col in view of Shang has taught a computer system as described in claim 14. Furthermore, it has been noted that the computer system of claim 15 performs the method of claim 8. Therefore, claim 15 is rejected for the same reasons set forth in the rejection of claim 8 above.

17. Referring to claim 19, Col in view of Shang has taught a method as described in claim 1.

Art Unit: 2183

- a) Col has further taught that the step of issuing comprises forcing the microinstructions to issue simultaneously, in lockstep with each other. See column 3, lines 52-56.
- b) Col has not explicitly taught that the step of canceling comprises canceling all of the plurality of microinstructions without regard to the relative ages of the microinstructions and without using a backoff mechanism. However, recall from the rejection of claim 1 above, that it would have been an obvious modification to add exception detection functionality to Col's system (in view of Shang). This modification, as described above, would allow for the detection of an exception within a single microinstruction and if an exception is triggered, the other related microinstructions will be subject to cancellation. Furthermore, recall that if interrupts are detected within Shang's system, results are not written to result registers, and consequently, there is no need for a backoff mechanism to undo results that have been incorrectly written.

18. Claim 5 is rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Shang, as applied above, and further in view of Hennessy and Patterson, Computer Architecture - A Quantitative Approach, 2nd Edition, 1996 (as applied in the previous Office Action, mailed on 12/20/2002, and herein referred to as Hennessy).

19. Referring to claim 5, Col in view of Shang has taught a method as described in claim 1. Col has further taught that his system can execute floating-point operations, which is supported in column 20, lines 32-38. Col has not explicitly taught that the microinstructions are executed over multiple clock cycles. However, it is well known and expected in the art that floating-point operations can consume more than one clock cycle for execution purposes. Hennessy has disclosed this concept on page 187. Hennessy has also shown a pipeline that accommodates

Art Unit: 2183

floating-point execution through multiple execution stages. See Fig.3.44 and Fig.3.45 on page 190. Since it has been disclosed by Hennessy that a floating-point operation takes multiple instruction execution cycles, it follows that it would have been obvious to one of ordinary skill in the art at the time of the invention to use a pipelined floating-point execution unit with multiple execution stages if floating-point operations are desired.

20. Claims 10, 11, 20, and 21 are rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Shang, as applied above, and further in view of Intel ®, Intel ® Architecture Optimization Reference Manual, 1998-1999 (herein referred to as Intel).

21. Referring to claim 10, Col has taught a method for processing software instructions comprising:

- a) providing two microinstructions to emulate a high-half and a low-half operation. See Fig.6 and note in cycle 7 that two microinstructions are executed in parallel. Col has not explicitly taught that the operation is an SSE operation. However, Intel has taught that SSE instructions are used to accelerate performance of applications regarding 3D geometry. See page 1-12. As a result, in order to increase system performance, it would have been obvious to one of ordinary skill in the art at the time of the invention to allow the system of Col to emulate SSE operations.
- b) forcing the high-half and low-half operations to issue in parallel. See column 3, lines 52-56.
- c) dispatching the high-half and low-half operations simultaneously to a first floating point unit and to a second floating point unit, respectively. See column 3, lines 52-56, and column 20, lines 32-38.

Art Unit: 2183

d) executing the high-half and low-half operations simultaneously, in lockstep. See column 3, lines 31-35, and Fig.6 (note in cycle 7 that two microinstructions are executed in parallel).

e) generating a signal from an emulator's hardware. Signals are inherently generated within a computer system. For instance, a clock signal is a basic signal that synchronizes the many different information-processing tasks assigned to the chip. Also, as instructions are fetched and decoded, signals are sent to the appropriate functional units in order to "specify" which operations are to be performed based on the type of instruction.

f) sending the signal to the first and second floating point functional units. Again, the appropriate signals would have to be supplied to the appropriate functional units in order to perform the desired operation.

g) Col has not explicitly taught determining whether an exception is taken in either the first or the second floating point unit and if the exception is taken in either the first or second floating point unit, preventing results from the high-half and low-half operations from being written to result registers, and canceling both the high-half and low-half operations. However, exceptions and the advantages of detecting exceptions are well known and expected in the art. In general, an exception is an interruption to the normal flow of program control, caused by the program itself or by executing an illegal instruction. Shang has taught a system in which macroinstructions are translated into a plurality of microinstructions and if an exception is detected within any one of those microinstructions, then the rest of the microinstructions are cancelled and results are not written to the result registers. Otherwise, if no exception is detected, then the results of the microinstructions are written to the result registers. See Fig.6 and note that if exceptions (interrupts) have been detected in at least one of the microinstructions

Art Unit: 2183

at step 104, then the system is flushed (cancellation of each microinstruction) at step 108. If no exceptions were detected at step 104, then the results are written to the result registers at step 106. In Col's system, since a macroinstruction is broken up into microinstructions, an exception in a single microinstruction would mean that an exception has occurred in the overall macroinstruction, and therefore, all of the microinstructions that represent a single macroinstruction, should not be able to change the state of the system. An advantage of Shang's scheme would be to avoid having to undo the changes made by the undesired execution of a microinstruction that is part of a faulty macroinstruction. This will prevent a reduction in throughput in that the extra time required to perform an undo-operation would not be necessary. Therefore, in order to maximize the efficiency of the overall system, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system such that it employs exception detection among microinstructions as taught by Shang.

h) Col has not explicitly taught updating MXCSR flags based upon the results of the first and second floating point units. However, the general idea of a status register is well known and expected in the art. The MXCSR register contains flags that are common to other processor status registers. These bits (flags) are set and cleared based on results from operations. For instance, if the result of an addition were zero, a flag indicating a zero-result would be set in the status register. Or, perhaps an overflow occurred. A flag in the status register would be set to specify that as well. Conditional statements such as branches then reference these flags in order to determine the direction of the program. Therefore, in order to provide a readable status of the processor so that programs (via branches) flow according to previously obtained results, it would

Art Unit: 2183

have been obvious to one of ordinary skill in the art at the time of the invention to update flags that are found in the MXCSR register.

22. Referring to claim 11, Col in view of Shang and further in view of Intel has taught a method as described in claim 10. Recall that it has been established that it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col such that it can detect and correct exceptions in a manner taught by Shang. Note also that Shang cancels all related microinstructions regardless of the relative ages of the microinstructions. See Fig.6.

23. Referring to claim 20, Col in view of Shang and further in view of Intel has taught a method as described in claim 10.

a) Col has further taught that the step of forcing the high-half and low-half operations to issue in parallel comprises causing the high-half and low-half operations to execute simultaneously in lockstep with each other. See column 3, lines 31-35, and Fig.6 (note in cycle 7 that two microinstructions are executed in parallel).

b) Col has not explicitly taught that flushing a result comprises canceling each of the high-half and low-half operations if an exception is taken in either the first or second floating point unit. However, as discussed in the rejection of claim 10 above, Shang has taught a system in which macroinstructions are translated into a plurality of microinstructions and if an exception is detected within any one of those microinstructions, then the rest of the microinstructions are cancelled and results are not written to the result registers. See Fig.6 and note that if exceptions (interrupts) have been detected in at least one of the microinstructions at step 104, then the system is flushed (cancellation of each microinstruction) at step 108. In Col's system, since a macroinstruction is broken up into microinstructions, an exception in a single microinstruction

Art Unit: 2183

would mean that an exception has occurred in the overall macroinstruction, and therefore, all of the microinstructions that represent a single macroinstruction, should not be able to change the state of the system. An advantage of Shang's scheme would be to avoid having to undo the changes made by the undesired execution of a microinstruction that is part of a faulty macroinstruction. This will prevent a reduction in throughput in that the extra time required to perform an undo-operation would not be necessary. Therefore, in order to maximize the efficiency of the overall system, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system such that it employs exception detection among microinstructions as taught by Shang.

24. Referring to claim 21, Col in view of Shang has taught a method as described in claim 1 wherein the step of executing comprises executing using a plurality of functional units of a floating-point unit. Col in view of Shang has not explicitly taught the emulation of SSE instructions. However, as discussed above, Intel has taught that SSE instructions are used to accelerate performance of applications regarding 3D geometry. See page 1-12. As a result, in order to increase system performance, it would have been obvious to one of ordinary skill in the art at the time of the invention to allow for the emulation of SSE instructions within the system of Col. Furthermore, it is inherent that within computer systems, as instructions are fetched and decoded, signals are sent to the appropriate functional units in order to "specify" which operations are to be performed based on the type of instruction. Therefore, if Col's system included SSE instructions, as established above, Col would have inherently taught:

Art Unit: 2183

a) generating a signal via hardware that indicates that the functional units are emulating an SSE instruction and sending the signal to the functional units. Again, the appropriate signals would have to be supplied to the appropriate functional units in order to perform the desired operation.

b) Also, it is inherent that determining an exception would occur after the functional unit has received its signal. For instance, the only way an overflow exception could be detected is by checking the result of an operation, which would only be obtained subsequent to “telling” the functional unit which operation to perform.

25. Referring to claim 22, Col in view of Shang has taught a system as described in claim 12. Furthermore, it has been noted that the system of claim 22 performs the method of claim 21. Therefore, claim 22 is rejected for the same reasons set forth in the rejection of claim 21.

26. Claim 9 is rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Shang, as applied above, and further in view of Phillips et al., U.S. Patent No. 6,038,652 (as applied in the previous Office Action, mailed on 12/20/2002, and herein referred to as Philips).

27. Referring to claim 9, Col in view of Shang has taught a method as described in claim 1.

a) Col has not explicitly taught that if an unmasked exception occurs, canceling the execution of all of the plurality of microinstructions, without regard to the relative ages of each of the plurality of microinstructions, and invoking a microcode handler. However, recall from the rejection of claim 1 above, that it would have been an obvious modification to add exception detection functionality to Col’s system (in view of Shang). This modification, as described above, would allow for the detection of an exception within a single microinstruction and if an exception is triggered, the other related microinstructions would be subject to cancellation. In

Art Unit: 2183

addition, Shang has taught that the triggering of an exception will result in the invocation of a microcode handler (referred to as an interrupt service routine). See column 11, lines 21-25. In general, the handler is invoked in order to correct the cause and effects of the exception and allow the processor to continue execution. Therefore, in order to correctly service an exception, it would have been obvious to one of ordinary skill in the art at the time of the invention to implement a microcode handler, which must be invoked upon exception detection.

b) Col in view of Shang has not explicitly taught updating at least one exception flag (when an unmasked exception occurs) by independently generating a logical OR of exceptions for a plurality of functional units. However, Phillips has taught the concept of simultaneously checking SIMD elements for exceptions and combining each individual exception into an overall exception. See FIG.2. Furthermore, the combining element (230) in FIG.2 can be implemented as an OR gate that generates a flag (240) used to specify whether or not an exception has occurred. See column 3, lines 60-63. A person of ordinary skill in the art would have recognized that the concept of Phillips would be applicable in Col's system in order to check for exceptions during the parallel execution of microinstructions. In a SIMD processor (as taught by Col), the overhead incurred to process the many possible exceptions generated by SIMD elements may be expensive and lead to degradation in performance. The system of Philips provides an efficient technique to report exceptions occurring in computing complex functions on a SIMD machine. An advantage to this scheme is that since an exception flag is produced according to a parallel execution of microinstructions as opposed to a serial execution of microinstructions, the processor will be able to detect an exception sooner and therefore sooner make the determination that the involved microinstructions should be cancelled. Therefore, it

Art Unit: 2183

would have been obvious to one of ordinary skill in the art at the time of the invention to update at least one exception flag in Col's system based on the exception check for a plurality of microinstructions.

28. Claim 18 is rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Shang and further in view of Intel, as applied above, and further in view of Makineni et al., U.S. Patent No 6,321,327 B1 (herein referred to as Makineni).

29. Referring to claim 18, Col in view of Shang and further in view of Intel has taught a computer system as described in claim 12. Col has further taught the general use of SIMD instructions, which is the format used by SSE instructions. See column 3, lines 61-63. However, Col has not disclosed the specific use of SSE instructions. Intel has taught that SSE instructions are used to accelerate performance of applications regarding 3D geometry. See page 1-12. As a result, in order to increase system performance, it would have been obvious to one of ordinary skill in the art at the time of the invention to implement an SSE instruction set within the system of Col. In addition, Col has taught that the SIMD execution units perform operations (such as an add) on multiple operands from a first SIMD register with corresponding multiple operands from a second SIMD register. See column 12, lines 1-8. Col has not explicitly taught the use of two 82-bit floating point registers for emulating four 32-bit single-precision floating-point values in an SSE register. However, Makineni has taught the use of 82-bit registers to hold 32-bit single precision floating-point numbers. See Fig.2B and Fig.3. Since SIMD operations involve performing a single operation on multiple pairs of data elements, it follows that multiple pairs of operands must be available to the SIMD execution unit. A person of ordinary skill in the art

Art Unit: 2183

would have recognized that by implementing 82-bit registers, multiple pairs of operands would be supplied to a SIMD execution unit by using multiple registers. In addition, by packing more than one data element into a single register, the amount of addressable registers could be decreased, resulting in less wires used for addressing purposes. Finally, each of the standard IEEE floating-point formats can be specified through the use of a single 82-bit floating point register, allowing the system to operate on different-precision operands depending on the situation. See the floating-point standards on page A-13 of Hennessy and note Fig.2A of Makineni shows a double-extended precision floating-point number. Therefore, in order to decrease the amount of hardware, while assuring the system has the capability of processing a wide variety of floating-point numbers, it would have been obvious to one of ordinary skill in the art to use two 82-bit floating point registers for emulating four 32-bit single-precision floating-point values in an SSE register.

Response to Arguments

30. Applicant's arguments filed on July 1, 2003, have been fully considered but they are not persuasive.

31. In the remarks, Applicant argues the novelty of claims 1, 10, and 12 on pages 6-7, in substance that:

“Applicants disclosed and incorporated by reference in a previously copending application...a mechanism for the forcing of the microinstructions to issue in parallel. Col allows for the parallel execution of instructions – Col does not require the forcing of the instructions.”

32. These arguments are not found persuasive for the following reasons:

Art Unit: 2183

a) Col has disclosed in column 3, lines 52-56, that two microinstructions are issued in parallel.

By issuing them in parallel, they are forced to issue in parallel.

33. In addition, Applicant argues on page 7, in substance that:

“...by forcing lockstep movement of the micro-operations, the claimed invention does away with the requirement for the extra storage that was required by the prior art. The combination of Col and Shang requires temporary storage, and thus does not teach how to overcome the problems solved by the claimed invention.”

34. These arguments are not found persuasive for the following reasons:

a) Shang has taught a system which detects an exception caused by one of a plurality of microinstructions which have been spawned via the translation of a macroinstruction. Shang has been combined with Col because Col has taught issuing and executing, in parallel, two microinstructions which have been spawned via the translation of a macroinstruction. As described in the previous action, implementing exception detection in Col's system would be beneficial so that a faulty microinstruction does not update the system. The fact that Shang uses temporary storage is irrelevant since the claims do not exclude the use of temporary storage when dealing with exception detection.

Conclusion

35. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire **THREE MONTHS** from the mailing date of this action. In the event a first reply is filed within **TWO MONTHS** of the mailing date of this final action and the advisory action is not mailed until after

Art Unit: 2183

the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to David J. Huisman whose telephone number is (703) 305-7811. The examiner can normally be reached on Monday-Friday (8:00-4:30).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Eddie Chan can be reached on (703) 305-9712. The fax phone numbers for the organization where this application or proceeding is assigned are (703) 746-7239 for regular communications and (703) 746-7238 for After Final communications.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.

DJH
David J. Huisman
August 12, 2003



EDDIE CHAN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100